IN THE KNOW

Proactive Defense: Threat Hunting & Deception

by James Motherway

Security, as we know it, is broken. According to IBM's "Cost of a Data Breach Report 2021,"[1] the average total cost of a breach has risen to US $4.24 million — the highest in history. Breaches are also taking longer to detect, with IBM reporting an average detection length of 287 days. Moreover, the impact of today's breaches is becoming more apparent and severe. Recently, in fact, in the midst of a pandemic, emboldened ransomware groups attacked critical infrastructure like hospitals. With advanced threats proliferating and new vulnerabilities emerging, company data appears to be an attacker's gold mine.

At the same time, many organizations fall into the trap of reacting to incidents as they occur, rather than proactively countering operational threats. They rely on basic security tools to mitigate adaptive and intelligent threats, leading to disastrous results. Some companies have serious detection blind spots; others collect more data than they can meaningfully use. Even as they lag attackers, security teams dream of someday getting ahead of them.

Fortunately, threat hunting, deception, and other proactive techniques are proving successful at disrupting and frustrating attackers. Indeed, some sophisticated organizations are finally able to realize the value from their security investments. But even as teams begin to put these practices to good use, leadership remains unsure of what a mature proactive security program might look like.

Certainly proactive techniques don't eliminate the need for basic security practices like those detailed in the Center for Internet Security (CIS) Critical Security Controls.[2] Asset inventories, audit log management, and strong incident-response capabilities are especially critical. Threat hunting requires both context and data, and organizations must be able to respond effectively when detecting an intrusion. Having said that, even organizations that aren't quite "there yet" can — and should — set threat hunting and deception as long-term goals. In this article, we describe advanced security operations of the type being used by major companies and the National Security Agency (NSA).

Reactive Security Isn't Enough

There are fundamental problems with the outmoded collect/watch/wait approach. At best, modern endpoint protection tools identify some attacker behavior at a generic level. At worst, signature-based defenses become outdated almost as soon as the signatures are written. Adding to the problem are untrained machine learning tools that create a deluge of false-positive alerts. These degrade the effectiveness of security teams and contribute to attrition. As defenders, we can't afford a minute lost to blind spots or distractions.

CrowdStrike's "2021 Global Threat Report" observed that interactive, human-driven intrusions have increased 400% in the past two years.[3] Most are from financially motivated criminal groups. A separate CrowdStrike report, "Nowhere to Hide," notes that these groups took about an hour and a half to move laterally across a network.[4] Just over a third took less than 30 minutes. Compromising your organization's security is their job, day in and day out.

As threats evolve, we can't rely on current security tools to do much more than detect yesterday's attacks. Breaking this cycle requires a strategic shift developed and carried out by informed, thoughtful practitioners. Proactive defense focuses our resources and maximizes their value, but it requires organizations to understand their threat landscape. When we know what we have and who might be after it, we can begin to picture how bad actors could accomplish their goals. Only then can we disrupt attacks earlier in their lifecycle, limit our damages, and possibly find enjoyment in doing so.

Threat-Hunting Theory

Threat hunting is an iterative, analytical process carried out by security practitioners. Parallels can be drawn with John Boyd's observe-orient-decide-act (OODA) loop. When engaged in an adversarial scenario, the person who navigates this cycle faster is typically victorious.

Threat hunting ideally reveals otherwise undetectable adversaries. However, poorly executed hunts run the risk of imparting a false sense of security. The threat-hunting process works best when supported by ample security data and well-defined hypotheses that can be proven (or disproven) by available information.

Here's an example:

> Reports from my sector state that threat actor Funny Cryptonym has used the Windows BITSAdmin tool while deploying malware. We're likely to be targeted based on our vertical, and we have endpoint and network data available to search. As a result, we should be able to identify whether we've also been affected by such a threat.

Faced with resistance, attackers will attempt to adopt unfamiliar techniques, slowing them down and making them easier to detect.

This hypothesis is informed by two important observations: threat intelligence and situational awareness. Hypotheses can also be produced via a practitioner's domain knowledge, such as past incident-handling experience.[5] In general, hunts are more successful when a number of elements are combined. However, finding intrusions is only one possible output of a hunt.

For example, if we identify that BITSAdmin is used legitimately in our network, the baseline knowledge of our environment will be improved. It's also possible to make separate observations, such as identifying misconfigurations or vulnerabilities. With that said, if this exercise did detect malicious activity, it would likely be ahead of the 287-day average detection time by an order of magnitude.

With this understanding, we can begin to conduct threat hunting in a proactive, structured manner. Threat hunters identify outliers in large data sets using techniques such as least-frequency analysis and time-series visualizations. They can identify patterns in disparate data sources that might otherwise go unmined.

This analysis informs subsequent hunts. Eventually, machine models can be trained based on feedback from threat hunters. Remaining tightly focused on an organization's unique concerns and requirements is the key to maximizing value.

Threat Hunting with MITRE ATT&CK

The MITRE ATT&CK knowledge base[6] is an excellent resource for threat hunting. ATT&CK decomposes and categorizes attacker behavior in detail. It can be used to understand the various points at which intrusions can be detected. These analyzed behaviors also have specific threat actors, malware, and mitigations mapped to them.

Before giving an example, some background information will be helpful. At a high level, the framework is composed of two key elements: tactics and techniques. Tactics correspond to the goals that an adversary must accomplish before completing an attack. Consider that attackers must first gain *initial access* to a network before deploying ransomware. Techniques are the means by which attackers achieve these goals.

Here we'll look at APT29 (also known as NOBELIUM), which was responsible for the SolarWinds supply chain attack.[7] Initial access was achieved through software supply chain compromise. However, the attackers also used relatively common techniques, including VBScripts to execute malware payloads and HTTP for C2 communication. Although supply chain compromise is difficult to mitigate, network intrusion prevention can help mitigate Web C2 traffic, as an example. (This is especially true when performing TLS interception.) Additionally, it's feasible for most companies to hunt for these more common techniques.

Hunting in this manner is incredibly powerful because it focuses on the behavior of attackers and their tools, rather than indicators like domain names or IP addresses. These are significantly more difficult and costly for attackers to change. Faced with resistance, attackers will attempt to adopt unfamiliar techniques, slowing them down and making them easier to detect. Other, more opportunistic, threat actors may simply move onto easier targets.

We can prioritize techniques to hunt on using visualization techniques. MITRE has released ATT&CK Navigator, a free open source tool to create and explore such visualizations (see Figure 1).[8] Recall that ATT&CK also has mitigations mapped to it. It's possible to create a heat map of control coverage, contrasted with the techniques of attackers in our threat model. This can be extremely beneficial in demonstrating to stakeholders that your program uses a methodical, threat-informed approach.

©2021 Cutter Consortium

Measuring Your Threat-Hunting Maturity Level

Hunting Maturity Model

Any discussion of threat hunting must include a mention of David Bianco's Hunting Maturity Model (HMM).[9] The model looks at an organization's level of security telemetry and the capability of its security analysts. It provides a way to gauge the level of a particular security program and where it could ultimately evolve to.

In Figure 2, HMM0 describes an organization with poor visibility that relies on out-of-the-box detection tools. No analysis is possible because little or no data is being collected, so reacting is the only possible action. At the next level (HMM1), the security function is more outward-looking. Analysts begin to consume threat intelligence and have a reasonable amount of data available to them, so threat hunting becomes possible. This is roughly the tipping point where a business moves away from a reactive security posture into something more valuable.

HMM2 is an ideal target for organizations wanting to maintain a threat-hunting program. Data is readily available for analysts, hunting occurs regularly using proven techniques, and valuable output can be expected. Finally, HMM3 and HMM4 describe organizations with the ability to use advanced or novel analysis techniques and increased automation. These are certainly excellent capabilities to have, but higher levels of maturity should be reserved for security teams with ample resources and a clear need for such investment.

Metrics

We need a clear way to demonstrate the value of threat hunting using metrics. It's important to note that during the inception of a threat-hunting program, there may be a spike in outputs followed by a leveling off. This is to be expected. The inverse can also occur later on, where relatively few findings culminate in a large spike as the function becomes more capable.

For organizations using the ATT&CK framework, hunts mapped to techniques can be an excellent metric.

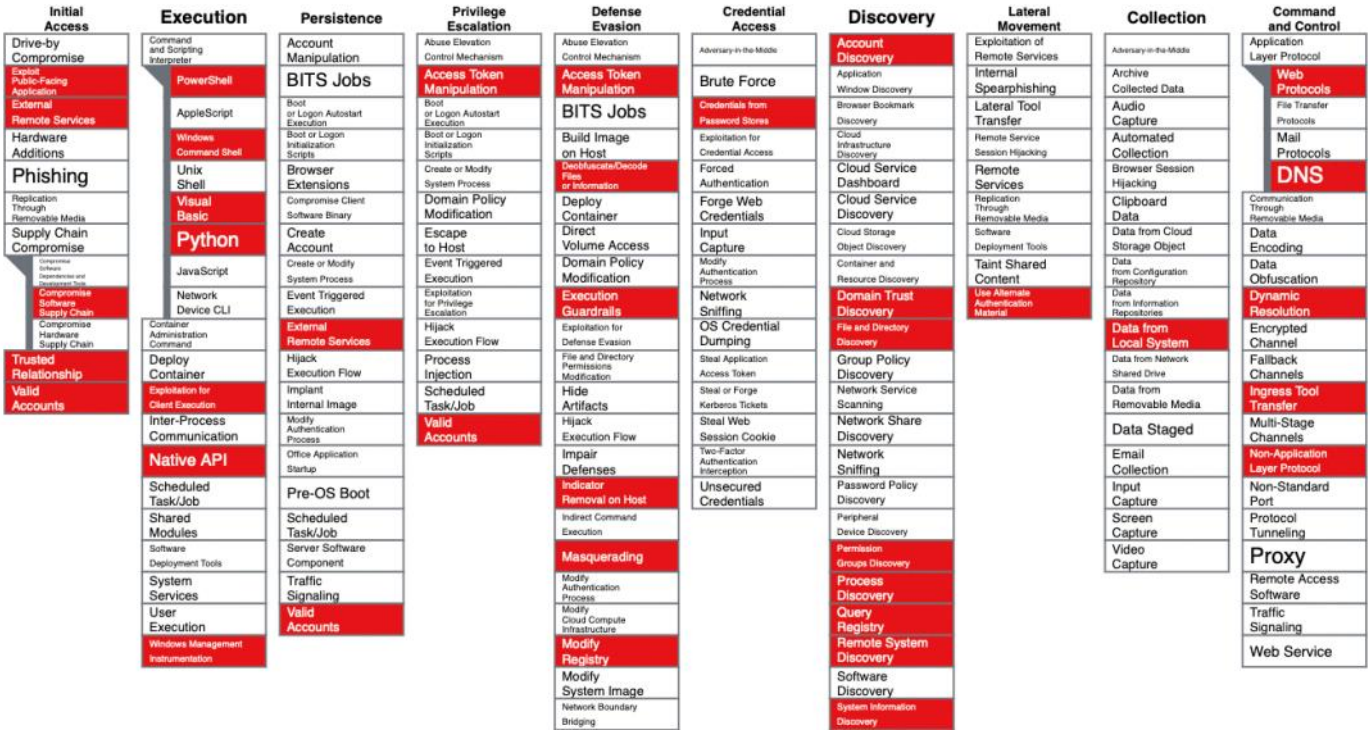

Figure 1 — APT29 SolarWinds techniques (including associated malware) from MITRE ATT&CK Navigator.

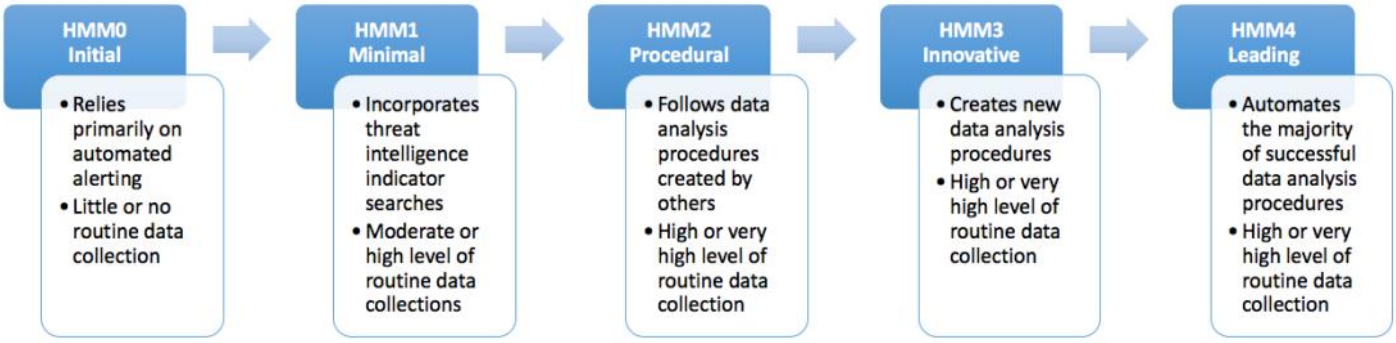
HMM0
Initial
• Relies primarily on automated alerting
• Little or no routine data collection
HMM1
Minimal
• Incorporates threat intelligence indicator searches
• Moderate or high level of routine data collections
HMM2
Procedural
• Follows data analysis procedures created by others
• High or very high level of routine data collection
HMM3
Innovative
• Creates new data analysis procedures
• High or very high level of routine data collection
HMM4
Leading
• Automates the majority of successful data analysis procedures
• High or very high level of routine data collection

Figure 2 — David Bianco's Hunting Maturity Model.

Techniques that have been prioritized using a method such as in Figure 1 can be tracked separately, as it demonstrates alignment with the organization's concerns. Similarly, it should be noted if mitigations are recommended as a result of identifying gaps.

It can also be useful to map which data sources have been used and/or in what environments hunts are carried out. Metrics can highlight biases that may indicate a lack of familiarity with log sources or a particular type of technology, such as the cloud. They can also demonstrate value from investments into tools that are heavily used. For organizations with less mature security programs, there could be issues with data collection or quality. When a hypothesis can't be tested fully due to this, it should be flagged and revisited later.

For organizations with robust incident-response capabilities, determining the time to detect a breach from initial compromise compared to industry average can underscore value. However, recall that malicious activity isn't the only outcome that we can measure; there are also:

- Newly identified vulnerabilities and/or misconfigurations
- Previously unknown architecture changes or shadow IT
- Hunts that have since been automated or refined

Using Deception

Deception isn't new in computer security; in fact, it dates back to at least the 1980s. One of the earliest accounts is Clifford Stoll's *The Cuckoo's Egg: Tracking a Spy Through the Maze of Computer Espionage.*[10] Stoll successfully used honeypot techniques to ensnare a nation-state threat actor. Other tools have since emerged, but few have proven as successful at degrading attacks.

In 2021, the National Institute of Standards and Technology (NIST) recognized deception as a notable tool in promoting resilience against advanced threats.[11] NSA's *The Next Wave* published an edition titled "Deception Research: Using Science to Reveal the Truth."[12] The article titled "Building the Science of Defensive Cyber Deception"[13] describes how NSA conducted a formal study to determine how deception could alter red teamers in a simulation:

> When unaware of the deception, the participants spent more time per host attempting to penetrate decoy assets than real ones. The decoy systems detected and alerted on all scanning and exploit activity related to the decoys. The participants verbally expressed confusion during the scenario and questioned their tools, their skills, and themselves rather than the authenticity of the network.

Surprisingly, this may be the least interesting finding of the study. Another cohort was merely told that deception was employed in the target environment, where in fact it wasn't. Still, the subjects progressed much slower and even exhibited paranoia to a degree. When encountering a genuinely vulnerable system, they suspected it was a trap.

A later, much larger, study measured how long it took for the first decoy to trigger during attack simulations.[14] Fifty percent were detected in 12 minutes, and 78% were detected in 20 minutes. A single outlier was detected in just under two minutes. This is in stark contrast with an industry average of 287 days and below the adversary breakout time of about 30 minutes. To date, there may be no other control that can boast such levels of efficacy.

©2021 Cutter Consortium

One reason deception is effective is that it exploits human vulnerabilities. Fear of detection, greed, and pressure to act quickly can all work in our favor as defenders. Unlike computers, the "software" on which we run on has a notoriously long patching cycle.

Deception in Practice

Attackers have resource constraints such as time and effort, and like all of us, tend to take the path of least resistance. Deception hinges on this. Its efficacy increases when this knowledge is fused with an understanding of adversary behavior. For organizations with threat-hunting capabilities, this could come naturally.

While *The Next Wave* edition mentioned previously focused on deception research primarily focused on "decoy assets" or perhaps low-interaction honeypots, there are ways to employ deception with minimal infrastructure. These can be explored with the help of Canarytokens, a tool that generates documents, links, and other assets that create an emailed alert when accessed.[15] The assets, generically referred to as "honeytokens," can inform defenders an intrusion may be in progress, and where to start looking. Those with file-monitoring capabilities can also identify what account made the access attempt, even in absence of the email functionality of the honeytoken.

However, those looking to adopt deception as a program of work may wish to look into managed solutions or running their own deployments. This can conceal the use of deception (avoiding references to Canarytokens infrastructure) and provides greater control. A combination of these could work just as well.

Files

Decoy Excel spreadsheets and PDFs are common honeytokens. They create detections by pulling online content when opened in their respective desktop applications. These are useful for flagging possible insider-threat behavior or threat actors that are less careful. If they're not opened in a particular manner, the detection won't trigger. Their coverage could be increased if benign documents are used, but these guide the reader to enticing locations using a link that actually refers to a DNS token.

Secrets and Keys

Credentials are like Christmas presents to criminals and penetration testers. As soon as they gain access to their target systems, they usually pillage the system for any passwords they can find. This helps intruder blend in by using valid accounts and gives them a way back in if evicted.

The best advice you can follow in this instance? Give the intruders what they want by hiding passwords in enticing documents that can be tucked away in email or storage. Passwords can even be injected into a system's memory, a common target for attackers looking to expand their repertoire of pilfered identities. In reality, those accounts are disabled, and the security team is closely monitoring for log-on attempts.

Attackers have resource constraints such as time and effort, and like all of us, tend to take the path of least resistance.

This logic can be applied to other types of secrets, such as keys. For example, Canarytokens provides a way to generate fake AWS keys that fire off an email with a predefined message when used. These can be left on devices or "leaked" in private containers and repositories. Creativity, methodical organization and thoughtful placement are essential to successful deception practices.

Servers

You can emulate servers and services for the sole purpose of confusing or enticing attackers. These low-interaction honeypots typically generate alerts when interacted with, providing limited or no functionality beyond that. When implemented correctly, they can confuse attackers about your network's topology and serve as enticing lures to attackers. These are more similar to the systems described in *The Next Wave*.

We can explore these decoys with the related OpenCanary project.[16] The daemon can run in a container or Linux system and can mimic databases and file shares, to name a few. Creating references to

Tread Lightly and Consult a Lawyer

When implementing a deception program, it's not enough to get buy-in from business and technology leaders. Speak with your general counsel about applicable laws and regulations (and remember that I'm a cybersecurity professional, not a lawyer).

In most cases, you'll need an airtight acceptable use policy. It should clearly state that corporate devices and the data on them belong to the company. It's also important to note that communications can and will be intercepted for purposes of securing the network. Upon accessing any device or system, a log-in banner should display that references the policy and says that continuing to the next step constitutes acceptance. The purpose is to retain prosecution as an option in the event of criminal activity.

The goal is to entice attackers while diminishing the likelihood that users will find them during the normal course of their duties. Ideally, it should take a deliberate, overt act to access the deception tool. When the deception creator fully understands his or her environment and collaborates fully with lines of business, the deception will create alerts with high true-positive rates early in the attack lifecycle.

these in your environment increases the likelihood of attracting (and detecting) any curious attackers who stumble across them.

When the techniques described here are combined, they have the potential to accelerate security programs that already have strong foundations in place.

It's Open Season

The old adage, "Defenders only need to make one mistake" is in need of review. I suggest: "Once on the network, attackers only need to make one mistake." Threat hunting and deception can be so effective that opportunistic criminals look elsewhere for their payday. Even nation-states may disengage when they encounter an environment that's actively hostile to intruders, fearing attribution for their activities. Then we might find ourselves closer to the day when defenders finally have the advantage.

References

[1]"Cost of a Data Breach Report 2021." IBM, 2021.

[2]"The 18 CIS Critical Security Controls." Center for Internet Security (CIS), 2021.

[3]"2021 Global Threat Report." CrowdStrike, 2021.

[4]"Nowhere to Hide: 2021 Threat Hunting Report." CrowdStrike, 2021.

[5]Lee, Robert M., and David Bianco. "Generating Hypotheses for Successful Threat Hunting." SANS Institute, 15 August 2016.

[6]MITRE ATT&CK knowledge base, accessed December 2021.

[7]"APT29." MITRE ATT&CK, accessed December 2021.

[8]"The ATT&CK Navigator: A New Open Source Project." MITRE, 5 March 2018.

[9]Bianco, David. "A Simple Hunting Maturity Model." Enterprise Detection and Response, 15 October 2015.

[10]Stoll, Clifford. *The Cuckoo's Egg: Tracking a Spy Through the Maze of Computer Espionage*. Gallery Books, 2005.

[11]Ross, Ron, et al. "Enhanced Security Requirements for Protecting Controlled Unclassified Information." National Institute of Standards and Technology (NIST), February 2021.

[12]Runser , Robert J., and Joe McCloskey (eds.). "Deception Research: Using Science to Reveal the Truth." *The Next Wave*, Vol. 23, No. 1, 2021.

[13]Rogers, Andrew, and Temmie Shade. "Building the Science of Defensive Cyber Deception." *The Next Wave*, Vol. 23, No. 1, 2021.

[14]Rogers and Shade (see 13).

[15]Canarytokens website, 2021.

[16]OpenCanary website, 2021.

James Motherway is a Senior Threat Hunt Analyst for National Grid. With expertise in incident response, secure configurations, and security awareness training, he helps organizations reach their desired target of cybersecurity maturity. Prior to his work at National Grid, Mr. Motherway developed a threat-hunting capability and has led security information and event management deployments and security awareness programs. Passionate about helping others, Mr. Motherway received a citation from US Governor of the State of Massachusetts Charlie Baker for a deployment with the American Red Cross. In his spare time, he enjoys building (and breaking) IT infrastructures. He can be reached at james@jamesmotherway.com.

©2021 Cutter Consortium